



[> home](#) [> about](#) [> feedback](#) [> log](#)

US Patent & Trademark Office



Try the *new* Portal  
design

Give us your opinion  
after using it.

## Search Results

Search Results for: `[(file <near/1> system)<AND>(((me <sentence>  
segment*) and (index* <sentence> file <sentence>  
segments))<AND>(meta_published_date <= 06-01-2001 )) )]`  
Found 19 of 121,005 searched.

Search within Results

[> Advanced Search](#) [> Search Help/Tips](#)

Sort by: [Title](#) [Publication](#) [Publication Date](#) [Score](#) [◆ Binder](#)

Results 1 - 19 of 19 [short listing](#)

1 [Performance aspects of disk space management](#) 100%

James O. Dyal , Michael K. Draughn

**Proceedings of the 13th conference on Winter simulation -  
Volume 1** January 1981

Dynamic Disk Space Management is a set of strategies for allocating files to disk storage devices so as to reduce seek time, device contention, and storage fragmentation. In this paper, the impact on system throughput and user response time is projected, via simulation model results, for a small general-purpose operating system which manages large capacity disk drives.

2 [Virtual memory management for database systems](#) 100%

Irving L. Traiger

**ACM SIGOPS Operating Systems Review** October 1982  
Volume 16 Issue 4

Over the last several years, a number of hardware and software systems have been developed which map entire files directly into the virtual memory address spaces used by programs. Since all file contents are directly addressable, there is no need for a programmer to issue explicit file system actions, such as Read or Write. In addition, all of the buffer management problems are eliminated, since programmers do not have to squeeze pieces of large files into small virtual spaces. Although these ad ...

- 3 Removing backing store administration from the CAP operating system 100%



Carl Dellar

**ACM SIGOPS Operating Systems Review** October 1980

Volume 14 Issue 4

- 4 Complex information processing: a file structure for the complex, the changing and the indeterminate 100%



T. H. Nelson

**Proceedings of the 1965 20th national conference** August 1965

THE KINDS OF FILE structures required if we are to use the computer for personal files and as an adjunct to creativity are wholly different in character from those customary in business and scientific data processing. They need to provide the capacity for intricate and idiosyncratic arrangements, total modifiability, undecided alternatives, and thorough internal documentation. I want to explain how some ideas developed and what they are. The original problem was to specify a comp ...

- 5 GLEAM/1130-a production system base for computer-aided design 100%




W. H. Sass

**Proceedings of the fifth annual 1968 design automation workshop on Design automation** July 1968

The partitioned computing system, operating several tasks concurrently, is an effective tool for computer-aided design. Such a system offers a high-performance man-machine interface and at the same time operates the low-speed graphic devices (such as plotters) continuously with small drain upon the system. The concept of

queueing tasks for devices, instead of queueing users for the system, improves user-access and maintains a reasonably constant load upon the system's hardware. Strong resid ...

**6** Query optimization I: Access paths in the "Abe" statistical query 100%


 facility

Anthony Klug

**Proceedings of the 1982 ACM SIGMOD international conference on Management of data** June 1982

An increasingly important part of information processing today involves the taking of counts, sums, averages, and other statistical or aggregate quantities. The "Abe" query language is designed to make formulation of complicated aggregations simple. Access path selection in Abe finds efficient ways to execute these complicated queries. Access paths for Abe queries perform "aggregate joins", that is, they compute aggregate quantities at the same time as they join subqueries with parent queries. T ...

**7** Modeling the storage architectures of commercial database systems 100%

 D. S. Batory

**ACM Transactions on Database Systems (TODS)** December 1985

Volume 10 Issue 4

Modeling the storage structures of a DBMS is a prerequisite to understanding and optimizing database performance. Previously, such modeling was very difficult because the fundamental role of conceptual-to-internal mappings in DBMS implementations went unrecognized. In this paper we present a model of physical databases, called the transformation model, that makes conceptual-to-internal mappings explicit. By exposing such mappings, we show that it is possible to model the storage ...

**8** Session II - recovery, concurrency and protection: Data base recovery 100%

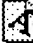
 at CMIC

Nicholas J. Giordano , Marvin S. Schwartz

**Proceedings of the 1976 ACM SIGMOD international conference on Management of data** June 1976

Several techniques have been combined to provide for data base recovery at CMIC. The CMIC environment is presented first in order to identify the constraints which data base recovery must satisfy. A technique is described for updating mass-storage structures (a B-tree in this case) in such a way that all information already addressable through the mass-storage structure can always be addressed through the mass-storage structure, even while the mass-storage structure is being updated. Audit trail ...


9 PDS/PIO: lightweight libraries for collective parallel I/O 100%

 Judy Sturtevant , Mark Christon , Philip D. Heermann , Pang-Chieh Chen

**Proceedings of the 1998 ACM/IEEE conference on Supercomputing (CDROM) November 1998**


PDS/PIO is a lightweight, parallel interface designed to support efficient transfers of massive, grid-based, simulation data among memory, disk, and tape subsystems. The higher-level PDS (Parallel Data Set) interface manages data with tensor and unstructured grid abstractions, while the lower-level PIO (Parallel Input/Output) interface accesses data arrays with arbitrary permutation, and provides communication and collective I/O operations. Higher-level data abstraction for finite element applic ...

10 Session 3: Objects and meta-objects: Matching data storage to application-needs 100%

 Dawson Dean , Richard Zippel

**Proceedings of the 6th workshop on ACM SIGOPS European workshop: Matching operating systems to application needs September 1994**


11 StorHouse metanoia - new applications for database, storage & data warehousing 100%

 Felipe Cariño , Pekka Kostamaa , Art Kaufmann , John Burgess  
**ACM SIGMOD Record , Proceedings of the 2001 ACM SIGMOD international conference on Management of data May 2001**  
Volume 30 Issue 2

This paper describes the StorHouse/Relational Manager (RM) database system that uses and exploits an *active storage hierarchy*. By active storage hierarchy, we mean that StorHouse/RM executes SQL queries *directly* against data stored on all hierarchical storage (i.e. disk, optical, and tape) without post processing a file or a DBA having to manage a data set. We describe and analyze StorHouse/RM features and internals. We also describe how StorHouse/RM differs from traditional HSM ...

## 12 Implementing NetUnify.com using squeak

100%

 Stephen Pair


**Addendum to the 2000 proceedings of the conference on Object-oriented programming, systems, languages, and applications (Addendum) January 2000**

NetUnify.com is a hosted service that enables its customers to maintain their project related material in a secure fashion over the Internet. They can grant access on a project-by-project basis to their employees, vendors, and customers. In addition, organizations can choose to operate their own NetUnify server rather than use the NetUnify.com service.

The NetUnify server is implemented in Squeak (see <http://www.squeak.org>), an open-source, highly portable object oriented environment. ...

## 13 The file system of an integrated local network

100%


 Paul J. Leach , Paul H. Levine , James A. Hamilton , Bernard L. Stumpf

**Proceedings of the 1985 ACM thirteenth annual conference on Computer Science March 1985**

The distributed file system component of the DOMAIN system is described. The DOMAIN system is an architecture for networks of personal workstations and servers which creates an integrated distributed computing environment. The distinctive features of the file system include: objects addressed by unique identifiers (UIDs);

transparent access to objects, regardless of their location in the network; the abstraction of a single level store for accessing all objects; and the layering of a network ...


**14** Personal student workstations: prospectus and requirements 100%

 Leslie J. Waguespack

**ACM SIGCSE Bulletin , Proceedings of the sixteenth SIGCSE technical symposium on Computer science education March 1985**

Volume 17 Issue 1

**15** A shared, segmented memory system for an object-oriented database 100%

 Mark F. Hornick , Stanley B. Zdonik

**ACM Transactions on Information Systems (TOIS) January 1987**

Volume 5 Issue 1

This paper describes the basic data model of an object-oriented database and the basic architecture of the system implementing it. In particular, a secondary storage segmentation scheme and a transaction-processing scheme are discussed. The segmentation scheme allows for arbitrary clustering of objects, including duplicates. The transaction scheme allows for many different sharing protocols ranging from those that enforce serializability to those that are nonserializable and require communi ...

**16** Serverless network file systems 100%

 Thomas E. Anderson , Michael D. Dahlin , Jeanna M. Neefe , David A. Patterson , Drew S. Roselli , Randolph Y. Wang


**ACM Transactions on Computer Systems (TOCS) February 1996**

Volume 14 Issue 1

We propose a new paradigm for network file system design: serverless network file systems. While traditional network file systems rely on a central server machine, a serverless system utilizes workstations cooperating as peers to provide all file system services. Any machine in the system can store, cache, or control any block of data. Our approach uses this location independence, in combination

with fast local area networks, to provide better performance and scalability th ...

**17** The memory extender personal filing system 100%


 W. P. Jones

**ACM SIGCHI Bulletin , Proceedings of the SIGCHI conference on Human factors in computing systems April 1986**

Volume 17 Issue 4

The benefits of electronic information storage are enormous and largely unrealized. As its cost continues to decline, the number of files in the average user's personal database may increase substantially. How is a user to keep track of several thousand, perhaps several hundred thousand, files? The Memory Extender (ME) system improves the user interface to a personal database by actively modeling the user's own memory for files and for the context in which these files are used. Files are mu ...

**18** A survey of power management techniques in mobile computing 100%

 operating systems

Gregory F. Welch

**ACM SIGOPS Operating Systems Review October 1995**

Volume 29 Issue 4

Many factors have contributed to the birth and continued growth of mobile computing, including recent advances in hardware and communications technology. With this new paradigm however come new challenges in computer operating systems development. These challenges include heretofore relatively unusual items such as frequent network disconnections, communications bandwidth limitations, resource restrictions, and power limitations. It is the last of these challenges that we shall explore in this p ...

**19** Energy efficient indexing on air 100%

 Tomasz Imielinski , S. Viswanathan , B. R. Badrinath

**ACM SIGMOD Record , Proceedings of the 1994 ACM SIGMOD international conference on Management of data May 1994**

Volume 23 Issue 2

We consider wireless broadcasting of data as a way of disseminating information to a massive number of users. Organizing and accessing information on wireless communication channels is different from the problem of organizing and accessing data on the disk. We describe two methods, (1,m) Indexing and Distributed Indexing, for organizing and accessing broadcast data. We demonstrate that the proposed algorithms lead to significant improvement ...

---

Results 1 - 19 of 19     [short listing](#)

---

The ACM Portal is published by the Association for Computing Machinery.  
Copyright © 2003 ACM, Inc.